

FAST '17: 15th USENIX Conference on File and Storage Technologies
February 27–March 2, 2017
Santa Clara, CA

Message from the Program Co-Chairs..... vii

Tuesday, February 28

Garbage

- Algorithms and Data Structures for Efficient Free Space Reclamation in WAFL**1
Ram Kesavan, Rohit Singh, and Travis Grusecki, *NetApp*; Yuvraj Patel, *University of Wisconsin–Madison*
- Tiny-Tail Flash: Near-Perfect Elimination of Garbage Collection Tail Latencies in NAND SSDs**15
Shiqin Yan, Huaicheng Li, Mingzhe Hao, and Michael Hao Tong, *University of Chicago*; Swaminathan Sundararaman, *Parallel Machines*; Andrew A. Chien and Haryadi S. Gunawi, *University of Chicago*
- The Logic of Physical Garbage Collection in Deduplicating Storage**29
Fred Douglass, Abhinav Duggal, Philip Shilane, and Tony Wong, *Dell EMC*; Shiqin Yan, *Dell EMC and University of Chicago*; Fabiano Botelho, *Rubrik, Inc.*

The System

- File Systems Fated for Senescence? Nonsense, Says Science!**45
Alex Conway and Ainesh Bakshi, *Rutgers University*; Yizheng Jiao and Yang Zhan, *The University of North Carolina at Chapel Hill*; Michael A. Bender, William Jannen, and Rob Johnson, *Stony Brook University*; Bradley C. Kuszmaul, *Oracle Corporation and Massachusetts Institute of Technology*; Donald E. Porter, *The University of North Carolina at Chapel Hill*; Jun Yuan, *Farmingdale State College of SUNY*; Martin Farach-Colton, *Rutgers University*
- To FUSE or Not to FUSE: Performance of User-Space File Systems**59
Bharath Kumar Reddy Vangoor, *Stony Brook University*; Vasily Tarasov, *IBM Research-Almaden*; Erez Zadok, *Stony Brook University*
- Knockoff: Cheap Versions in the Cloud**73
Xianzheng Dou, Peter M. Chen, and Jason Flinn, *University of Michigan*
- HopsFS: Scaling Hierarchical File System Metadata Using NewSQL Databases**89
Salman Niazi, Mahmoud Ismail, Seif Haridi, and Jim Dowling, *KTH Royal Institute of Technology*; Steffen Grohsschmiedt, *Spotify AB*; Mikael Ronström, *Oracle*

Edward Sharpe and the Magnetic Zeros

- Evolving Ext4 for Shingled Disks**105
Abutalib Aghayev, *Carnegie Mellon University*; Theodore Ts'o, *Google, Inc.*; Garth Gibson, *Carnegie Mellon University*; Peter Desnoyers, *Northeastern University*
- SMaRT: An Approach to Shingled Magnetic Recording Translation**121
Weiping He and David H.C. Du, *University of Minnesota*
- Facilitating Magnetic Recording Technology Scaling for Data Center Hard Disk Drives through Filesystem-Level Transparent Local Erasure Coding**135
Yin Li and Hao Wang, *Rensselaer Polytechnic Institute*; Xuebin Zhang, *Dell EMC/DSSD*; Ning Zheng, *Rensselaer Polytechnic Institute*; Shafa Dahandeh, *Western Digital*; Tong Zhang, *Rensselaer Polytechnic Institute*

(Continues on next page)

Wednesday, March 1

Corruption

- Redundancy Does Not Imply Fault Tolerance: Analysis of Distributed Storage Reactions to Single Errors and Corruptions**149
Aishwarya Ganesan, Ramnatthan Alagappan, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*
- Omid, Reloaded: Scalable and Highly-Available Transaction Processing**167
Ohad Shacham, *Yahoo Research*; Francisco Perez-Sorrosal, *Yahoo*; Edward Bortnikov and Eshcar Hillel, *Yahoo Research*; Idit Keidar, *Technion—Israel Institute of Technology and Yahoo Research*; Ivan Kelly, *Midokura*; Matthieu Morel, *Skyscanner*; Sameer Paranjpye, *Arimo*
- Application Crash Consistency and Performance with CCFS**.....181
Thanumalayan Sankaranarayanan Pillai, Ramnatthan Alagappan, and Lanyue Lu, *University of Wisconsin—Madison*; Vijay Chidambaram, *The University of Texas at Austin*; Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*
- High Performance Metadata Integrity Protection in the WAFL Copy-on-Write File System**197
Harendra Kumar; Yuvraj Patel, *University of Wisconsin—Madison*; Ram Kesavan and Sumith Makam, *NetApp*

Frameworks

- Mirador: An Active Control Plane for Datacenter Storage**213
Jake Wires and Andrew Warfield, *Coho Data*
- Chronix: Long Term Storage and Retrieval Technology for Anomaly Detection in Operational Data**229
Florian Lautenschlager, *QAware GmbH*; Michael Philippsen and Andreas Kumlehn, *Friedrich-Alexander-Universität Erlangen-Nürnberg*; Josef Adersberger, *QAware GmbH*
- Crystal: Software-Defined Storage for Multi-Tenant Object Stores**243
Raúl Gracia-Tinedo, Josep Sampé, Edgar Zamora, Marc Sánchez-Artigas, and Pedro García-López, *Universitat Rovira i Virgili*; Yosef Moatti and Eran Rom, *IBM Research—Haifa*

Solid State Records

- WORT: Write Optimal Radix Tree for Persistent Memory Storage Systems**257
Se Kwon Lee, *UNIST (Ulsan National Institute of Science and Technology)*; K. Hyun Lim, *Hongik University*; Hyunsub Song, Beomseok Nam, and Sam H. Noh, *UNIST (Ulsan National Institute of Science and Technology)*
- SHRD: Improving Spatial Locality in Flash Storage Accesses by Sequentializing in Host and Randomizing in Device**.....271
Hyukjoong Kim and Dongkun Shin, *Sungkyunkwan University*; Yun Ho Jeong and Kyung Ho Kim, *Samsung Electronics*
- Graphene: Fine-Grained IO Management for Graph Computing**285
Hang Liu and H. Howie Huang, *The George Washington University*

Thursday, March 2

Faster Faster

vNFS: Maximizing NFS Performance with Compounds and Vectorized I/O301

Ming Chen, *Stony Brook University*; Dean Hildebrand, *IBM Research-Almaden*; Henry Nelson, *Ward Melville High School*; Jasmit Saluja, Ashok Sankar Harihara Subramony, and Erez Zadok, *Stony Brook University*

On the Accuracy and Scalability of Intensive I/O Workload Replay315

Alireza Haghdoost and Weiping He, *University of Minnesota*; Jerry Fredin, *NetApp*; David H.C. Du, *University of Minnesota*

On the Performance Variation in Modern Storage Stacks329

Zhen Cao, *Stony Brook University*; Vasily Tarasov, *IBM Research-Almaden*; Hari Prasath Raman, *Stony Brook University*; Dean Hildebrand, *IBM Research-Almaden*; Erez Zadok, *Stony Brook University*

Enlightening the I/O Path: A Holistic Approach for Application Performance.345

Sangwook Kim, *Apposha and Sungkyunkwan University*; Hwanju Kim, *Sungkyunkwan University and Dell EMC*; Joonwon Lee and Jinkyu Jeong, *Sungkyunkwan University*

Open Channel D

LightNVM: The Linux Open-Channel SSD Subsystem359

Matias Bjørling, *CNEX Labs, Inc. and IT University of Copenhagen*; Javier Gonzalez, *CNEX Labs, Inc.*; Philippe Bonnet, *IT University of Copenhagen*

FlashBlox: Achieving Both Performance Isolation and Uniform Lifetime for Virtualized SSDs375

Jian Huang, *Georgia Institute of Technology*; Anirudh Badam, Laura Caulfield, Suman Nath, Sudipta Sengupta, and Bikash Sharma, *Microsoft*; Moinuddin K. Qureshi, *Georgia Institute of Technology*

DIDACache: A Deep Integration of Device and Application for Flash Based Key-Value Caching391

Zhaoyan Shen, *Hong Kong Polytechnic University*; Feng Chen and Yichen Jia, *Louisiana State University*; Zili Shao, *Hong Kong Polytechnic University*